

Efficient Action Detection in Untrimmed Videos via Multi-Task Learning

Yi Zhu and Shawn Newsam
University of California, Merced
{yzhu25, snewsam}@ucmerced.edu

Abstract

This paper studies the joint learning of action recognition and temporal localization in long, untrimmed videos. We employ a multi-task learning framework that performs the three highly related steps of action proposal, action recognition, and action localization refinement in parallel instead of the standard sequential pipeline that performs the steps in order. We develop a novel temporal actionness regression module that estimates what proportion of a clip contains action. We use it for temporal localization but it could have other applications like video retrieval, surveillance, summarization, etc. We also introduce random shear augmentation during training to simulate viewpoint change. We evaluate our framework on three popular video benchmarks. Results demonstrate that our joint model is efficient in terms of storage and computation in that we do not need to compute and cache dense trajectory features, and that it trains several times faster than its sequential ConvNets counterpart. Yet, despite being more efficient, it outperforms state-of-the-art methods with respect to accuracy.

1. Introduction

The enormous explosion of user-generated videos on the Internet represents a wealth of information. It is essential to understand and analyze these videos for tasks like video recommendation and search, video highlighting, video surveillance, human-robot interaction, human skill evaluation, etc. Automatically localizing and recognizing human actions or events in long, untrimmed videos can save tremendous manual effort and cost. Much progress has been made on action detection in manually trimmed, short videos over the past several years [47, 14, 13, 30, 39, 6] by employing hand-crafted local features, such as improved dense trajectories (IDT) [35], or video representations that are learned directly from the data itself using deep convolutional neural networks (ConvNets). However, these approaches are much less effective when applied to long, untrimmed, realistic videos due to the amount of background clutter (non-action clips) as well as the large variation in scene context, motion

speed, action style, and camera viewpoint.

Presently, the dominant approach to action detection [10, 12, 43, 28] is a sequential pipeline usually consisting of three steps, either using IDT features plus shallow classifiers or through deep neural networks. First, a video is processed in a temporal sliding window manner to generate action proposals. The candidate proposals are then used to train a more sophisticated action classifier for recognition. Finally, the temporal boundaries of action instances are estimated through post-processing by applying duration priors or non-maximum suppression (NMS). While this pipeline has achieved state-of-the-art performance in recent years [23, 36, 16, 45], we argue that *learning and applying the closely related steps independently has several drawbacks*. First, the performance is upper-bounded by the proposal stage. Second, sequential training prevents the modules from sharing information and restricts the training data each stage sees. Together, these limit the generalizability of the model especially when there is large intra-class variation. Finally, training the stages sequentially is computationally expensive especially for large-scale video analysis.

Inspired by multi-task learning, such as the work on facial landmark detection in [46], we develop a joint learning framework based on 3D ConvNets [32]. Our goal is to answer the following three questions simultaneously: (1) *Does a video clip contain an action of interest?* (2) *What action class does it belong to?* and (3) *How much temporal actionness does it have?* These three problems are closely related and so a joint learning scheme should have better regularization. This will improve the network’s capacity to discriminate and make it converge faster compared to its sequential ConvNets counterpart [28].

Our multi-task learning framework for action detection is outlined in Figure 1. The top “Proposal” branch is a binary classifier that determines whether a video clip is an action of interest or just background. The middle “Classification” branch is a conventional action recognition module that classifies a video clip into different action categories. This branch utilizes multiple intermediate losses calculated from last convolutional layer (conv5) as well as the first (fc6) and last (fc8) fully connected layers in order to im-

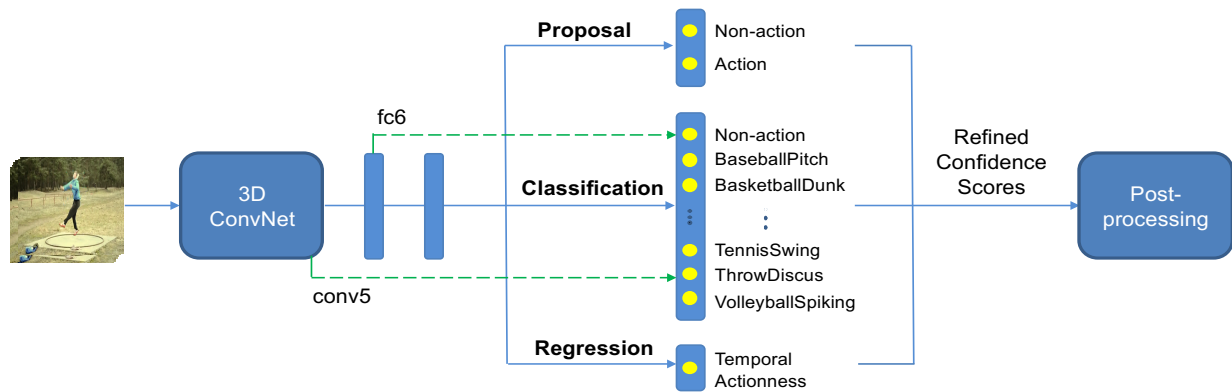


Figure 1: Overview of our multi-task learning framework for action detection. The 3D ConvNet is initialized from [32]. The three branches of our model simultaneously perform complementary tasks: *Proposal* to determine if a video clip is an action or background, *Classification* of a clip into categories, and *Regression* of what proportion of a clip contains action. The two green dashed lines indicate the multiple intermediate losses used in the action classification branch. The individual outputs are combined to produce refined prediction scores which are used in post-processing to obtain the final detection results.

prove regularization. The bottom “Regression” branch predicts what proportion of a video clip contains the action of interest temporally. This temporal actionness is related to but different from the spatial actionness introduced in [1]. That quantity is purely spatial and refers instead to the likelihood that a specific image region contains a generic action instance.

We show that our joint learning scheme can perform action detection using a single 3D ConvNet without needing to compute dense trajectory features. Additionally, we introduce a data augmentation technique based on shear transformation to increase view-point invariance. This augmentation is shown to be particularly effective for action recognition in untrimmed videos. We also include multiple intermediate losses during training to help prevent overfitting. Our multi-task network is significantly more efficient to train, requiring only a fraction of the time of a cascaded ConvNets approach, yet achieves better performance on three popular benchmarks, ranging from THUMOS 2014 (complex) [15] to MSR Actions II (simple) [44] and MPII cooking (fine-grained) [27].

2. Related Work

There exists an extensive body of literature on video analysis and human action detection. Here we review only the most related work.

Action Localization: This field has been studied from two directions: spatio-temporal action localization and temporal action localization.

Spatio-temporal action localization is the more active research area and aims to localize an action simultaneously

in space and time. There is a sizable body of work on this problem [20, 14, 13, 30, 39, 6, 19]; however, most of this work focuses on short, trimmed videos. In addition, spatio-temporal localization requires substantial effort to manually generate the object bounding boxes needed for training. This is not feasible for large-scale video analysis such as unconstrained Internet videos.

In contrast, our focus is on temporal action localization. Given an untrimmed video, we want to determine the temporal bounds of one or more action instances. Many real world applications only need the start and end times of the action/event in a video stream. Further, temporal action localization can serve as a relatively fast pre-processing step that allows more expensive spatio-temporal localization methods to be applied to just those time periods in which action is detected. Yeung et al. [42] employ a recurrent neural network (RNN)-based agent and use REINFORCE [40] to learn the agent’s decision policy. Their fully end-to-end approach achieves promising results on two large-scale benchmarks while only needing to observe a fraction of the video frames. Shou et al. [28] adapt 3D ConvNets [32] for action recognition and design a new multi-stage framework. They explicitly take temporal overlap into consideration which benefits the final post-processing step to obtain state-of-the-art results on THUMOS 2014. Sun et al. [31] use noisy tagged web images to discover localized action frames in videos and model temporal information with long short-term memory (LSTM) networks [11]. Very recently, Heilbron et al. [10] introduced a sparse learning framework to represent and retrieve activity proposals efficiently. While our approach shares some structural similarity with these works, it is distinct in that it exploits regu-

larities between closely related tasks for action detection in a multi-task learning framework.

Multi-task Learning: Multi-task deep networks have become a popular approach to solve multiple semantically related tasks in computer vision. In general, they can be formulated as a fan-out model, where the earlier layers share a single representation and the later layers diverge. Multi-task learning has shown promising results in various domains, such as action parsing [20], semantic segmentation [2], natural language processing [21], object detection [5], etc.

Among these works, the DAP3D-Net in [20] is most closely related to us. They use a general fan-out model to automatically parse actions in videos and output action types, action attributes, and action locations, simultaneously. However, our work is different from theirs in several key ways: (a) [20] performs spatio-temporal action localization, while our goal is temporal action localization; and (b) they train on well-aligned synthetic action data. In contrast, we utilize widely adopted, untrimmed, realistic video benchmarks. We take inspiration from [28] in designing our multi-task learning framework. But turning a cascaded pipeline into a single network is not easy. In particular, care must be taken to address the data imbalance problem when training a parallel network. We also introduce a training data augmentation technique based on shear transformation, design a multiple intermediate losses fusion strategy, and develop a temporal actionness regression module. These contributions highlight the novelty of our joint learning framework compared to previous work on action detection in large-scale untrimmed videos.

3. Overview

Our goal is to develop a model that can determine the temporal bounds of one or more action instances in a long, untrimmed video. We particularly want the model to be efficient during training and testing (application) in terms of both computation and storage.

Problem Formulation: For an untrimmed video dataset \mathbb{V} , each video V may contain one or multiple action instances \mathbb{A} . Here, V consists of a set of individual frames v_t , where t indicates the time index of the frame. During training, we generate fixed-length video clips from the annotated videos via temporal sliding windows. These clips are the inputs to our multi-task learning framework. The network is trained using five losses: a softmax loss of action/non-action, three softmax losses of action categorization, and a Euclidean squared loss of temporal actionness. During prediction, a video is decomposed into clips which are fed to the trained model to obtain per-clip action categorization scores from the last fully connected layer. These scores are weighted by the output of the other two network branches: the proposal confidence score and the estimated temporal actionness. Finally, non-maximum suppression is applied using the re-

finer action categorization scores, and the top- k results are selected as the temporally localized actions of interest.

Motivation: There are several reasons to develop a multi-task learning framework. First, most existing work does not fully exploit the usefulness of background videos to help train a more robust classifier. Most background videos are discarded after the action proposal stage in order to produce a cleaner and more balanced training data set. Further, as shown in [38], the overall action recognition/detection performance depends on the performance of the action proposal module which is never perfect and will incorrectly label some true actions as background. In addition, the prediction scores of a video clip from the classification module can be high even if it only has a few action frames and the rest are background. This can be problematic because subsequent post-processing steps, such as non-maximum suppression, might remove clips with overall lower action categorization scores but which temporally have a large proportion of action [28]. We therefore develop a temporal actionness module to quantitatively measure the proportion of action in a video clip. In sum, closely related tasks can be integrated in a multi-task learning framework without a reduction in accuracy but with much faster training and/or prediction performance [20, 2, 21, 5]. Our three main objectives are *efficiency*, *accuracy* and *robustness*.

4. Methodology

In this section, we first describe our training data augmentation techniques. We then describe our network architecture and the new regression module to predict temporal actionness. Finally, we provide details on post-processing using the refined action categorization scores.

4.1. Training Data Augmentation

Most video benchmarks for action detection are extremely small compared to the large-scale image datasets like ImageNet [3]. For example, the largest video dataset with temporal annotations to-date, ActivityNet [9], only contains 15,410 training action instances. This makes overfitting a challenge in deep learning. Data augmentation techniques, such as random flipping [18], corner cropping [37], multi-scale spatial cropping [7, 29], and multi-scale temporal sliding windows [28], have therefore been devised to improve the training of deep ConvNets.

In this paper, we utilize three standard data augmentation techniques and introduce a new one for action detection. We apply multi-scale temporal sliding windows to generate the training video clips. For an untrimmed video, we extract temporal sliding windows of lengths 16, 32, 64, 128, 256, and 512 frames [28]. For window lengths 16 and 32, the sliding stride is 16, while for other lengths, it is 32. This strategy effectively suppresses a bias for shorter windows during NMS [22]. For each window, we uniformly sample




Figure 2: Illustration of our random shear training data augmentation. Top: Sample video frames of action “High jump” from the THUMOS challenge 2014 training set. Bottom: We pick the middle image from the top row and shear it between -25 and 25 degrees for display. While subtle, the difference between the leftmost and rightmost frames does simulate limited viewpoint change and results in better models. When training, we use just one sheared sample (with a random amount of shear) to augment the data. This figure is best viewed in color.

16 frames to form our fixed-size input clips. All clips are spatially resized to have a frame size of 128×171 pixels [32]. During training, we horizontally flip each input video clip with 50% probability. We also do corner-cropping, extracting the four corners and the centers of the frames.

We introduce random shear transformation to simulate viewpoint change for action detection problems. This is motivated by the fact that the same action can appear quite different depending on the viewpoint. For example, sample video frames of the action “High jump” are shown in the top row of Figure 2. The intra-class variability due to viewpoint change is quite large. The ideal augmentation technique would generate multi-view training clips from a single-view input. This would help reduce overfitting. However, generating 3D videos from 2D is a difficult and computationally expensive problem itself [24, 4, 41], and cannot be done on-the-fly while training deep networks.

We therefore instead use shear transformations to efficiently simulate limited viewpoint changes. We perform random shearing with a shear parameter θ between -25 and 25 degrees. This range is determined using cross-validation (too small a value makes no difference while too large introduces unrealistic distortions). Border elements are repeated for padding and the sheared image is cropped to its original size. Sample sheared frames can be seen in the bottom row of Figure 2. While subtle, shear transformations are shown to simulate limited viewpoint change. Compare the leftmost and rightmost frames. In the experiments below, we show quantitatively that this augmentation strategy does help prevent overfitting and improves the final detection accuracy on all three benchmarks.

4.2. Network Architecture

In [32], the authors show that 2D ConvNets “forget” the temporal information in the input signal after every convolution operation. They propose 3D ConvNets which an-

alyze sets of contiguous video frames organized as clips. 3D ConvNets have shown to be effective at learning spatio-temporal features in video volume data analysis problems [33] such as video understanding, optical flow estimation, video coloring, etc.

We therefore base our framework on the 3D ConvNets model C3D described in [32]. Our network has 8 convolutional layers, 5 pooling layers, followed by two fully connected layers, and three parallel outputs. All 3D convolution filters are $3 \times 3 \times 3$ with stride 1, and all 3D pooling layers are $2 \times 2 \times 2$ with stride 1, except for `pool1`. In order to preserve temporal information in the earlier layers, `pool1` has size $1 \times 2 \times 2$. The detailed network structure is shown in Figure 1. Since C3D is trained on the large Sports1M video dataset [17], it serves as a good initialization point for training our model. The inputs to our network are minibatches of video clips of size $16 \times 3 \times 112 \times 112$, and the outputs are three sets of confidence scores: action or not, action category, and temporal actionness.

In order to reduce overfitting and increase the influence of the action categorization branch, we introduce two additional loss functions. Besides the softmax loss from layer `fc8`, we add two more softmax losses to regularize the learning, one from layer `conv5` and another from layer `fc6`. The benefits of this are two-fold. (i) Layer `conv5` contains weak semantic, but strong spatial information, which is important for motion understanding; layers `fc6` and `fc8` can provide high-level semantic and abstract information about the video clip. These three loss functions thus incorporate complementary information from the top layers of the network. (ii). Incorporating these additional intermediate losses promotes action categorization as the most important of the three main tasks. Improved action recognition also benefits the other two branches. Training our network thus considers a total of five losses.

Note that, though there is strong motivation to perform temporal action localization via a parallel multi-task framework, the joint learning of the tasks is a challenge due to potential imbalance in the training data. The training data that is optimal for one branch of the network might not be optimal for the others. The action proposal branch is a binary classification problem and seeks to distinguish between two classes, action and non-action. In contrast, the action categorization branch is a multi-class classification problem and seeks to distinguish between $N + 1$ classes (N classes plus background). Usually $N > 1$ so if we generate balanced training data for the action proposal branch, it will be unbalanced for the action categorization branch, and vice versa. We explore two ways to address this problem. (i) Provide balanced data for the action proposal task and tune the class weights when calculating the action categorization loss. This essentially means giving more attention to action clips when training the categorization branch. Or, (ii) pro-

vide balanced data for the action categorization branch and tune the class weights when calculating the action proposal loss. We show in Section 5 that (ii) results in higher accuracy. One possible reason is that action categorization is our main goal, while action proposal is just a supporting task. If we provide balanced data for the action proposal task, our network sees too many non-action examples and learns features more from background video clips than from actions of interest.

4.3. Temporal Actionness

Actionness was introduced in [1] to indicate the likelihood that a specific spatial location contains a generic action instance. This measure is useful for generating accurate action proposals, as well as for related tasks such as action recognition/detection. It applies to the spatial dimension only though and is thus more like a spatial saliency measure. It works well for spatio-temporal action localization, but is not appropriate for temporal action localization.

Hence, we introduce a new concept, *temporal actionness*, that is better suited to volume data analysis, especially when using 3D ConvNets. Temporal actionness is a fractional value p_a , ranging from 0 to 1, that indicates what proportion of a video volume contains an action of interest temporally. During training, p_a is computed as the temporal intersection of a video clip with the actual action period, divided by the clip’s length, and is thus a scalar. When p_a equals 0, the volume contains no action. When p_a equals 1, all the frames contain the action. If p_a is limited to two values, 0 or 1, this problem reduces the standard binary classification problem of action or non-action.

Compared to the temporal Intersection over Union (tIoU) metric, temporal actionness is more effective because its computation does not depend on the length of ground truth, which means this quantity can be predictable. It based only on how much action of interest a clip contains. On the contrary, predicting the value of tIoU is basically impossible when the ground truth information is missing. This may be the primary reason that direct learning of a regressor for time shift, video segment duration, or tIoU usually does not work well in practice.

We design a regression module as shown in Figure 1 to learn temporal actionness. To demonstrate its effectiveness, we compare it with another approach from [28]. That method contains a localization stage with a new loss function and uses tIoU to scale the confidence scores generated by the classification stage. It achieves state-of-the-art performance when training the modules separately. To compare approaches, we simply replace our regression module with their loss function. The results in Section 5 show that our proposed model achieves better performance. The benefit of our approach is that an explicit prediction of temporal actionness could also have other applications.

4.4. Post-processing

For evaluation, given a video V , we use multi-scale temporal sliding windows to generate input video clips $\{v_m\}_{m=1}^M$ where m indicates the m -th clip and M is the total number of clips for V . For each v_m , we obtain three outputs: the non-action confidence score p_b^m , the action categorization confidence score \mathbf{p}_1^m , and the estimated temporal actionness p_a^m . Bold font indicates vectors.

We design a hierarchical method to integrate these three results. *Clip-level*: For each v_m , we refine \mathbf{p}_1^m by multiplying by the temporal actionness value p_a^m . This boosts the categorization confidence scores in clips with lots of action. *Video-level*: For the whole video V , we refine \mathbf{p}_1^m through:

$$w_m \times \mathbf{p}_1^m, \quad w_m = \frac{e^{-\alpha p_b^m}}{\sum_n e^{-\alpha p_b^n}} \quad (1)$$

where w_m is the softmax weights used to scale the categorization confidence scores of the m -th video clip [38]. In our pipeline, video clips that are predicted as background are not discarded, they are just suppressed. The lower p_b^m , the higher w_m , which means the prediction scores of a video clip will be increased if it is more likely to be an action clip, and they will be decreased if it is more likely to be a non-action clip. α is a parameter which can be tuned to find the best weighting.

After the clip-level and video-level refinement, we re-score each window using the frequency of its duration in the training set as a per-class duration prior [22]. However, this re-scoring step may be optional because no substantial improvement was observed. Note that we can reduce inference time by discarding non-action clips (i.e. when p_b^m is smaller than some threshold) as most existing works do. Finally, we apply non-maximum suppression with overlap threshold δ , and select the top- k segments to form the final detection result.

5. Experiments and Results

In this section, we first describe the implementation details and the benchmark datasets used for evaluation. We then present the experimental results and study the effects of the various design choices described in section 4.

5.1. Implementation Details

During training, the network weights are learned using mini-batch (30 video clips) stochastic gradient descent with momentum (set to 0.9). The learning rate is initialized to 0.0001, except for layer **fc8** with 0.01. Both learning rates decrease to 1/10 of their values whenever the performance saturates, and training is stopped when the learning rate is smaller than 10^{-6} . Dropout is applied with a ratio of 0.5. The five losses are fused with equal weights which has shown the best performance through empirical validation.

Model	$\alpha = 0.1$	$\alpha = 0.2$	$\alpha = 0.3$	$\alpha = 0.4$	$\alpha = 0.5$
Karaman et al. [16]	4.6	3.4	2.1	1.4	0.9
Wang et al. [36]	18.2	17.0	14.0	11.7	8.3
Oneata et al. [23]	36.6	33.6	27.0	20.8	14.4
Sun et al. [31]	12.4	11.0	8.5	5.2	4.4
Heilbron et al. [10]	36.1	32.9	25.7	18.2	13.5
Richard et al. [26]	39.7	35.7	30.0	23.2	15.2
Yeung et al. [42]	48.9	44.0	36.0	26.4	17.1
Ours fc8	45.6	41.2	34.5	26.1	17.0
Ours conv5 + fc8	46.6	42.9	35.6	28.5	18.7
Ours conv5 + fc6 + fc8	47.7	43.6	36.2	28.9	19.0

Table 1: Action detection results on THUMOS 2014 with varied IoU threshold α . All performances are reported using mAP. Top section: Top three performers in the 2014 challenge. Middle section: Recent state-of-the-art approaches. Bottom section: Our approach. We observe that our network with multiple intermediate losses performs the best. Note that conv5, fc6, fc8 here refer to losses, not features.

For a temporal localization evaluation metric, we follow the standard protocol in THUMOS challenge 2014 to calculate the mean average precision (mAP). A detection is considered to be correct if its overlap with the ground truth action instance is above some threshold, and the action category is predicted correctly.

5.2. THUMOS Challenge 2014

THUMOS challenge 2014 is considered to be one of the most challenging datasets for temporal action detection. This dataset consists of 13,320 trimmed videos from 101 actions for training, 1,010 temporally untrimmed videos for validation, and 1,574 untrimmed videos for testing. The detection task involves 20 out of the 101 classes, which are all sports related. During testing, we only use 213 test videos with annotations to evaluate our framework.

Results: Since our goal is temporal action localization for untrimmed videos, we include the validation videos in our training data. We have already described how we generate temporal multi-scale video clips for untrimmed videos in Section 4. For trimmed videos, we uniformly sample 16 frames as clips without multi-scale augmentation, and each clip has the same action label as the video. During training, we perform 15K iterations with learning rate 10^{-4} , 15K iterations with 10^{-5} , and 10K iterations with 10^{-6} . During prediction, the overlap threshold in non-maximum suppression is set to 0.4.

Table 1 shows that our method outperforms previous state-of-the-art approaches. The three models in top section of Table 1 use Fisher vector encoded IDT features on video clips and perform post-processing to refine the localization scores. [23] is the winning approach of the 2014 challenge. It uses video-level action classification scores as a context feature which greatly improves the performance. The four methods in middle section of Table 1 are the most recent state-of-the-art results. Note that [42] achieves bet-

Training Set	$\alpha = 0.2$	$\alpha = 0.5$
$V_T + V_U$	43.5	19.0
$\frac{3}{4}V_T + V_U$	41.7	17.9
$\frac{1}{2}V_T + V_U$	36.9	14.4

(a) Sequential Network [28]

Training Set	$\alpha = 0.2$	$\alpha = 0.5$
$V_T + V_U$	43.6	19.2
$\frac{3}{4}V_T + V_U$	42.9	18.7
$\frac{1}{2}V_T + V_U$	39.4	17.3

(b) Our Parallel Network

Table 2: Comparison between our parallel multi-task learning pipeline and its sequential counterpart in [28]. Our method is robust to large reductions in training data. Note the mAP here is reported using the updated evaluation toolkit.

ter accuracy in the small α regime. One possible reason is that their learned observation policy tends to output shorter segments compared to a standard sliding window approach. Significantly, though, we estimate that the computation time required to train the approach of [42] is much higher than ours. Their observation network, a fine-tuned VGG-16 model, by itself requires approximately the same amount of computation time to train as our single 3D ConvNets. They also have to train the RNN-based agent to learn the decision policy. [28] reports a mAP of 19.0 when $\alpha = 0.5$. This is using the updated evaluation toolkit [15], however, and so is not directly comparable with the results in Table 1. For completeness, we also evaluate our result using the updated toolkit and obtain an accuracy of 19.2 given our multi-loss fusion network. Furthermore, our approach is almost 3 times faster to train than [28] and we only need to store one model instead of three. More importantly, our method is simpler to train than the multi-stage pipeline of [28] which requires manually balanced datasets and separate input files for the three stages.

Figure 3 shows three examples of predicted action instances from the THUMOS 2014 challenge test set. Our approach is shown to accurately localize the action instances.

5.3. Discussion

In this subsection, we will discuss the benefits of using our parallel learning framework over its sequential counterpart. We choose THUMOS 2014 as the illustrating dataset. **Enhanced Regularization** Multi-task learning can benefit from the enhanced regularization provided by closely-related tasks. This can increase a model’s robustness, especially under fewer training samples regime. Here, we explore this advantage of our joint learning scheme by gradually reducing the number of trimmed videos in the training set. Our training set for THUMOS 2014 is composed of two parts, the original trimmed training videos which we denote as V_T , and the untrimmed validation videos which we denote as V_U . By decreasing the proportion of V_T in our training set, we can validate our framework’s robustness.

As shown in Table 2(b), we achieve the best performance using the full training dataset, as expected. Notably, though, the performance decreases only slightly when just 3/4 of




Figure 3: Examples of predicted action instances on THUMOS 2014. Faded frames indicate background. Ground truth is green while our predictions are blue. Top: Cliff-Diving. Middle: LongJump. Bottom: ThrowDiscus. This figure is best viewed in color.

V_T is used. The relative reduction is still only 9.8% when $1/2$ of V_T is used.

By contrast, less training data is much more likely to result in overfitting in a multi-stage approach like [28]. To demonstrate this, we trained the second “action categorization” stage of [28] using $1/2$ of V_T and all V_U and the classification accuracy of the 3D ConvNets dropped dramatically. This, in turn, will degrade the performance of the whole system since the localization network is initialized from the categorization stage. The final action detection mAP is 14.4 as shown in Table 2(a). The relative performance decrease is 24.2% for the sequential framework. This demonstrates that our joint learning framework is more robust when fewer trimmed video examples are available. This is important for video analysis because current video datasets do not have large numbers of labeled training samples.

Computational and Storage Efficiency Our approach is computationally efficient compared to approaches that compute and fuse expensive features [23] or to approaches with multiple stages [28, 10, 42]. Our single model only requires 40K iterations to train and each iteration takes approximately one second. Although non-action video clips are not discarded, our class-balanced sampling strategy effectively controls how many clips our ConvNet sees during training. For THUMOS 2014, our model sees less than 10% of the video frames. For the evaluation, our network is two times faster than the multi-stage pipeline [28]. We can further increase test time efficiency by controlling the proportion of non-action video clips to preserve by setting a threshold.

With regards to storage, our approach does not need to store high-dimensional representations, such as Fisher Vector encoded features. And, only a single ConvNet model is required for deployment whose size is less than 0.35GB.

Ablation Study We conduct ablative experiments to understand the effects of the different components and parameters in our framework.

First, we compare different strategies for balancing the

Data Balance	0.2	0.5
Proposal	40.1	17.9
Categorization	43.6	19.0

(a)

Shear Param	0.2	0.5
$\theta = 0$	40.6	18.3
$\theta = 10$	40.9	18.3
$\theta = 25$	43.6	19.0
$\theta = 40$	39.7	17.8

(c)

Overlap	0.2	0.5
Localization loss [28]	43.2	18.9
Temporal Actionness	43.6	19.0

(b)

Contribution	0.2	0.5
w/o proposal	41.7	18.2
w/o regression	40.0	17.5
full model	43.6	19.0

(d)

Table 3: Action detection results of our framework on THUMOS 2014 under different configurations. (a) Data imbalance during training. (b) Effectiveness of proposed temporal actionness. (c) Random shear augmentation. (d) Impact of each branch. See the text in Ablation Study for more details.

training data. As shown in Table 3(a), balancing for the action categorization task yields better performance, especially when $\alpha = 0.2$. As mentioned before, balancing for the action proposal task likely results in the network seeing too many negative examples and features that are tuned more to background video clips than to actions of interest.

In Table 3(b), we replace our regression module based on temporal actionness with the well-designed localization loss function proposed in [28]. Our approach obtains slightly better performance, but the temporal actionness explicitly predicted by our network can have other applications.

We also determine the optimal value of the shear parameter θ from among 10, 25 and 40 degrees. In Table 3(c), $\theta = 0$ means there is no shear augmentation. When $\theta = 10$, the performance does not change. One possible reason is that the corner-cropping augmentation might already provide similar benefits. When $\theta = 40$, the performance decreases by 0.5%, possibly due to too much distortion. For $\theta = 25$, we observe an improvement of 0.7% when $\alpha = 0.5$. Considering that the best mAP reported is 19.0 when tIoU is 0.5, an improvement from random shear of 0.7 is relatively significant. Most importantly, random shear improves the accuracy of the “classification” branch by 1.6 percent on THUMOS 2014. After the post-processing and non-maximum suppression, the improvement reduces to 0.7 for the final action “detection”. Hence, random shear augmentation is very useful to help overcome overfitting during model training.

Finally, we explore whether our action proposal and temporal actionness branches are redundant since the former can be seen as an extreme case of the latter. As shown in Table 3(d), the full model achieves the best performance, indicating the two branches are complementary. The *w/o proposal* model performs 0.8% worse than the full model, while the *w/o regression* model performs 1.5% worse. This demonstrates that estimating temporal actionness is more beneficial than performing action proposal.

Model	$\alpha = 0.5$	Model	$\alpha = 0.5$
Yu et al. [43]	28.2	Sliding Window	7.9
Gemert et al. [34]	54.5	Gemert et al. [34]	13.1
Heilbron et al. [10]	60.3	Richard et al. [26]	14.0
Ours + fc8	59.6	Ours + fc8	14.6
Ours conv5 + fc6 + fc8	61.1	Ours conv5 + fc6 + fc8	14.9

(a) MSR Action II

(b) MPII Cooking

Table 4: Action detection results on MSR Action Dataset II and MPII Cooking dataset. Our approach consistently achieves better performance compared to state-of-the-art approaches at 0.5 overlap threshold.

5.4. MSR Action Dataset II

The MSR Action Dataset II [44] consists of 54 video sequences (203 action instances) recorded in a crowded environment. There are three action types: hand waving, hand-clapping, and boxing. We follow the standard cross-dataset evaluation protocol of using the KTH dataset for training. During training, we perform 8K iterations with learning rate 10^{-4} , 8K iterations with 10^{-5} , and 4K iterations with 10^{-6} . During prediction, the overlap threshold in non-maximum suppression is set to 0.5.

Table 4(a) shows that our joint learning framework with multiple intermediate losses performs better than the state-of-the-art approach of [10]. That method uses IDT features to represent the video clips, however, which prevents it from being applied to large-scale datasets since the Fisher vector encoded features need to be cached. Our approach can easily scale to large datasets. We attribute our good performance to the joint regularized learning and to the temporal actionness refined post-processing. Note that, [43] and [34] are designed to generate spatio-temporal proposals. [10] projects these proposals to the temporal dimension and randomly picks a subset of temporal proposals to perform action detection. An example of predicted “boxing” instances using our approach can be seen in Figure 4 top.

5.5. MPII Cooking

MPII Cooking [27] is a large, fine-grained cooking activities dataset. It contains 44 videos with a total length of more than 8 hours of 12 participants performing 65 different cooking activities. It consists of a total of 5,609 annotations spread over the 65 activity categories, including a background class for the action detection task. It also has a second type of annotation, articulated human pose for a pose estimation problem, which we ignore. Following the standard protocol in [27], we have 7 splits after performing leave-one-person-out cross-validation. Each split use 11 subjects for training, leaving one for validation. During training, we perform 10K iterations with learning rate 10^{-4} , 10K iterations with 10^{-5} , and 10K iterations with 10^{-6} . During prediction, the overlap threshold in non-maximum suppression is set to 0.5.



Figure 4: Examples of predicted action instances. Faded frames indicate background. Ground truth is green while our predictions are blue. Top: “Boxing” in MSR Action Dataset II. Bottom: “Take out from drawer” in MPII Cooking dataset. This figure is best viewed in color.

We compare our method to a sliding window baseline similar to [26, 27]. Table 4(b) shows our method performs better than both the baseline and recent state-of-the-art approaches [26, 34]. [26] is novel in that it includes a length and language model in addition to an action classifier. However, their runtime during inference is quadratic in the number of frames. By limiting the maximal action length to a constant, they can solve the action detection problem in a reasonable time, but this not easily scalable to long videos. When we compare our approach to [34], we report their performance by applying the implementation kindly provided by the authors. An example of predicted “Take out from drawer” instances using our approach can be seen in Figure 4 bottom.

6. Conclusion

We present a novel multi-task learning framework to detect human actions in long, untrimmed videos. Performing the three related tasks in parallel exploits the training data more thoroughly and is more efficient in time and space. The random shear augmentation simulates viewpoint differences. The multiple intermediate losses force the network to learn better representations for action categorization. These theoretical insights are supported by our experimental results on three popular datasets. We will make our trained models publicly available for other researchers. In future work, we plan to revise our multi-task approach to use a 3D residual learning framework [8]. We are also interested in optimizing the network end-to-end similarly to Faster R-CNN [25], since our temporal actionness scores can be considered as region proposal scores.

Acknowledgements. This work was funded in part by a National Science Foundation CAREER grant, #IIS-1150115, and a seed grant from the Center for Information Technology in the Interest of Society (CITRIS). We gratefully acknowledge the support of NVIDIA Corporation through the donation of the Titan X GPU used in this work.

References

- [1] W. Chen, C. Xiong, R. Xu, and J. J. Corso. Actionness Ranking with Lattice Conditional Ordinal Random Fields. In *CVPR*, 2014.
- [2] J. Dai, K. He, and J. Sun. Instance-aware Semantic Segmentation via Multi-task Network Cascades. In *CVPR*, 2016.
- [3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009.
- [4] J. Flynn, I. Neulander, J. Philbin, and N. Snavely. DeepStereo: Learning to Predict New Views from the World’s Imagery. *arXiv preprint arXiv:1506.06825*, 2015.
- [5] M. Ghifary, W. B. Kleijn, M. Zhang, and D. Balduzzi. Domain Generalization for Object Recognition with Multi-task Autoencoders. In *ICCV*, 2015.
- [6] G. Gkioxari and J. Malik. Finding Action Tubes. In *CVPR*, 2015.
- [7] K. He, X. Zhang, S. Ren, and J. Sun. Spatial Pyramid Pooling in Deep Convolutional Networks for Visual Recognition. *TPAMI*, 2015.
- [8] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016.
- [9] F. C. Heilbron, V. Escorcia, B. Ghanem, and J. C. Niebles. ActivityNet: A Large-Scale Video Benchmark for Human Activity Understanding. In *CVPR*, 2015.
- [10] F. C. Heilbron, J. C. Niebles, and B. Ghanem. Fast Temporal Activity Proposals for Efficient Detection of Human Actions in Untrimmed Videos. In *CVPR*, 2016.
- [11] S. Hochreiter and J. Schmidhuber. Long Short-Term Memory. *Neural Computation*, 1997.
- [12] M. Jain, J. van Gemert, H. Jgou, P. Boutheymy, and C. G. M. Snoek. Action Localization with Tubelets from Motion. In *CVPR*, 2014.
- [13] M. Jain, J. C. van Gemert, T. Mensink, and C. G. M. Snoek. Objects2action: Classifying and Localizing Actions without Any Video Example. In *ICCV*, 2015.
- [14] M. Jain, J. C. van Gemert, and C. G. M. Snoek. What do 15,000 Object Categories Tell Us about Classifying and Localizing Actions? In *CVPR*, 2015.
- [15] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. 2014.
- [16] S. Karaman, L. Seidenari, and A. D. Bimbo. Fast Saliency Based Pooling of Fisher Encoded Dense Trajectories, 2014.
- [17] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale Video Classification with Convolutional Neural Networks. In *CVPR*, 2014.
- [18] A. Krizhevsky, I. Sutskever, and G. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *NIPS*, 2012.
- [19] T. Lan, Y. Wang, , and G. Mori. Discriminative Figure-Centric Models for Joint Action Localization and Recognition. In *ICCV*, 2001.
- [20] L. Liu, Y. Zhou, and L. Shao. DAP3D-Net: Where, What and How Actions Occur in Videos? *arXiv preprint arXiv:1602.03346*, 2016.
- [21] M.-T. Luong, Q. V. Le, I. Sutskever, O. Vinyals, and L. Kaiser. Multi-task Sequence to Sequence Learning. In *ICLR*, 2016.
- [22] D. Oneata, J. Verbeek, and C. Schmid. Action and Event Recognition with Fisher Vectors on a Compact Feature Set. In *ICCV*, 2013.
- [23] D. Oneata, J. Verbeek, and C. Schmid. The LEAR Submission at THUMOS 2014, 2014.
- [24] V. Ramakrishna, T. Kanade, and Y. Sheikh. Reconstructing 3D Human Pose from 2D Image Landmarks. In *ECCV*, 2012.
- [25] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *NIPS*, 2015.
- [26] A. Richard and J. Gall. Temporal Action Detection using a Statistical Language Model. In *CVPR*, 2016.
- [27] M. Rohrbach, S. Amin, M. Andriluka, and B. Schiele. A Database for Fine Grained Activity Detection of Cooking Activities. In *CVPR*, 2012.
- [28] Z. Shou, D. Wang, and S.-F. Chang. Temporal Action Localization in Untrimmed Videos via Multi-stage CNNs. In *CVPR*, 2016.
- [29] K. Simonyan and A. Zisserman. Two-Stream Convolutional Networks for Action Recognition in Videos. *NIPS*, 2014.
- [30] K. Soomro, H. Idrees, and M. Shah. Action Localization in Videos through Context Walk. In *ICCV*, 2015.
- [31] C. Sun, S. Shetty, R. Sukthankar, and R. Nevatia. Temporal Localization of Fine-Grained Actions in Videos by Domain Transfer from Web Images. In *ACM MM*, 2015.
- [32] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning Spatiotemporal Features with 3D Convolutional Networks. In *ICCV*, 2015.
- [33] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Deep End2End Voxel2Voxel Prediction. In *CVPRW*, 2016.
- [34] J. van Gemert, M. Jain, E. Gati, and C. G. Snoek. APT: Action Localization Proposals from Dense Trajectories. In *BMVC*, 2015.
- [35] H. Wang and C. Schmid. Action Recognition with Improved Trajectories. In *ICCV*, 2013.
- [36] L. Wang, Y. Qiao, and X. Tang. Action Recognition and Detection by Combining Motion and Appearance Features. 2014.
- [37] L. Wang, Y. Xiong, Z. Wang, and Y. Qiao. Towards Good Practices for Very Deep Two-Stream ConvNets. *arXiv preprint arXiv:1507.02159*, 2015.
- [38] Y. Wang and M. Hoai. Improving Human Action Recognition by Non-action Classification. In *CVPR*, 2016.
- [39] P. Weinzaepfel, Z. Harchaoui, and C. Schmid. Learning to Track for Spatio-Temporal Action Localization. In *ICCV*, 2015.
- [40] R. J. Williams. Simple Statistical Gradient-following Algorithms for Connectionist Reinforcement Learning. *Machine Learning*, 1992.
- [41] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3D ShapeNets: A Deep Representation for Volumetric Shapes. In *CVPR*, 2015.
- [42] S. Yeung, O. Russakovsky, G. Mori, and L. Fei-Fei. End-to-end Learning of Action Detection from Frame Glimpses in Videos. In *CVPR*, 2016.

- [43] G. Yu and J. Yuan. Fast Action Proposals for Human Action Detection and Search. In *CVPR*, 2015.
- [44] J. Yuan, Z. Liu, and Y. Wu. Discriminative Video Pattern Search for Efficient Action Detection. *TPAMI*, 2011.
- [45] J. Yuan, Y. Pei, B. Ni, P. Moulin, and A. Kassim. ADSC Submission at THUMOS Challenge 2015. 2015.
- [46] Z. Zhang, P. Luo, C. C. Loy, and X. Tang. Facial Landmark Detection by Deep Multi-task Learning. In *ECCV*, 2014.
- [47] Y. Zhu and S. Newsam. Depth2Action: Exploring Embedded Depth for Large-Scale Action Recognition. In *ECCVW*, 2016.